that data packets are only transmitted over communication links having transmission costs that are lower than their variable translucency TTL value.

Accordingly, by decrementing the variable translucency TTL value of a data packet by the transmission cost associated with the communication link over which the packet is to be transmitted, a routing protocol may be adapted for more efficient use in connection with networks composed at least in part of high-capacity links. Variable translucency TTLs thus enable improved routing accuracy and expanded horizon lines in networks capable of tolerating the increased link-state overhead associated with this increased accuracy, without significantly increasing the link-state overhead in networks incapable of tolerating such increases in such overhead. The unacceptable trade-off discussed above in connection with routing protocols employing conventional TTLs (where relatively minor reductions in link-state traffic overhead are achieved in consideration for reduced routing accuracy) is thus overcome, resulting in increased network efficiency and operability.

Although the above description of variable translucency TTLs has been provided in connection with a no-sight routing scheme, persons of skill in the art will recognize that variable translucency TTLs are generally applicable to any number of routing schemes employing data packets in which the distance information may propagate is limited. For example, in addition to no-sight routing schemes, variable translucency TTLs might also be adapted for use in connection with hazy-sighted routing schemes or traditional link-state routing schemes in which the life of a packet is limited by a TTL.

II. Exemplary Process Flows

a. No-Sight Routing

FIG. **4** depicts an exemplary process flow **400** for performing no-sight routing. Process flow **400** is meant to serve merely as an example for how no-sight routing may be implemented, such that the various steps recited therein should not be deemed to limit the scope of its application. In addition, various other intermediary steps not explicitly illustrated in FIG. **4** may be performed as needed, regardless of their exclusion from FIG. **4**.

At step **402**, the various parameters of each mobile node are configured. In many embodiments, these parameters are configured by system administrators of the network and/or device manufacturers of the mobile nodes. According to certain embodiments, a propagation limit and/or an algorithm or heuristic for computing a propagation limit, as discussed above, is stored in the mobile node during step **402**. As described in greater detail above, a mobile node may be assigned one or more propagation limits during this step for use in connection with all types of data to be transmitted by the mobile node. A mobile node may also be assigned a different propagation limit for each data packet type to be transmitted by this mobile node. In addition, although all mobile nodes within a network may be configured to have the same propagation limit, the propagation limits of each mobile node may be modified as needed.

Step **402** may occur in an automated fashion, such as by a computer, so that manual configuration is unnecessary. In addition, step **402** may be omitted from exemplary process flow **400** as needed, such as when propagation limits are to be dynamically assigned or adjusted based on network properties such as node density, as discussed in greater detail above.

Upon completion of step **402**, control proceeds to step **404** where a mobile node transmits an LSP to all neighboring nodes. As described above, link-state packets typically contain data detailing the ID of the node that created the LSP and a list of directly connected neighbors of that node. In many embodiments, the travel of this LSP throughout the network is determined, as explained in greater detail above, by a propagation limit.

At step **406**, LSPs from neighboring nodes are received. Again, due to the propagation limits placed on these LSPs, not all nodes connected to the network will receive all other nodes' LSPs. Upon receipt of these LSPs from neighboring nodes, at step **408** each node computes a limited network map based on the information received from each LSP received. As various topology changes to the network are perceived by a mobile node (i.e., a neighboring mobile node is added or removed), the mobile node generates and transmits to its neighboring devices LSUs detailing the changes that have been observed.

b. Variable Translucency TTLs

FIG. **5** depicts an exemplary process flow **500** for propagating data based on variable translucency time-to-live values. Process flow **500** is meant to serve merely as an example for how variable translucency TTLs may be utilized, such that the various steps recited therein should not be deemed to limit the scope of its application. In addition, various other intermediary steps not explicitly illustrated in FIG. **5** may be performed as needed, regardless of their exclusion from FIG. **5**.

At step **502**, a source node in exemplary network **300** determines to which of its neighboring nodes, if any, a data packet is to be transmitted. In many embodiments, the source node makes this determination in accordance with procedures provided by a routing protocol, such as the no-sight routing protocol described above.

Once a destination node has been chosen, at step **504** the transmission cost of the communication link provided between the source node and this destination node is determined. In certain embodiments, transmission costs are assigned to communication links by a routing protocol. The transmission cost of a communication link may also be computed by a source node in accordance with, for example, a predetermined heuristic. For example, a source may compute the transmission cost of a communication link based on various parameters of the link, such as the physical length of the link, the bandwidth capacity of the link, the type of medium used to facilitate the link, etc.

Upon determining the transmission cost of the communication link connecting the source node to the destination node, at step **506** it is determined whether the variable translucency TTL value of the data packet to be transmitted is greater than the transmission cost of the communication link. If so, control proceeds to step **508**. Otherwise, control returns to step **502** where another destination node, if available, is chosen.

At step **508**, the variable translucency TTL value of the data packet to be transmitted is decremented by the determined/computed transmission cost of the communication link. At step **510**, the source node transmits the data packet to the chosen destination node.

III. Alternative Embodiments

In accordance with the provisions of the patent statutes, the principles and modes of operation have been explained and illustrated. However, it should be understood that embodiments described herein may be practiced otherwise than is specifically explained and illustrated without departing from the spirit or scope thereof, and that the invention described herein is intended to be limited only by the following claims.